arXiv:1511.08350v1 [cs.AI] 26 Nov 2015

A global Constraint for mining Sequential Patterns with GAP constraint

Amina Kemmar[1], Samir Loudni[2], Yahia Lebbah[1]
Patrice Boizumault[2], and Thierry Charnois[3]

[1]LITIO – University of Oran 1 – Algeria, EPSECG of Oran – Algeria
[2] GREYC (CNRS UMR 6072) – University of Caen – France
[3]LIPN (CNRS UMR 7030) – University PARIS 13 – France

Abstract. Sequential pattern mining (SPM) under gap constraint is a challenging task. Many efficient specialized methods have been developed but they are all suffering from a lack of genericity. The Constraint Programming (CP) approaches are not so effective because of the size of their encodings. In [7], we have proposed the global constraint PREFIX-PROJECTION for SPM which remedies to this drawback. However, this global constraint cannot be directly extended to support gap constraint. In this paper, we propose the global constraint GAP-SEQ enabling to handle SPM with or without gap constraint. GAP-SEQ relies on the principle of right pattern extensions. Experiments show that our approach clearly outperforms both CP approaches and the state-of-the-art `cSpade` method on large datasets.

1 Introduction

Mining sequential patterns (SPM) is an important task in data mining. There are many useful applications, including discovering changes in customer behaviors, detecting intrusion from web logs and finding relevant genes from DNA sequences. In recent years many studies have focused on SPM with gap constraints [16,18]. Limited gaps allow a mining process to bear a certain degree of flexibility among correlated pattern elements in the original sequences. For example, [6] analyses purchase behaviors to reflect products usually bought by customers at regular time intervals according to time gaps. In computational biology, the gap constraint helps discover periodic patterns with significant biological and medical values [14].

Mining sequential patterns under gap constraint (GSPM) is a challenging task, since the *apriori property* does not hold for this problem: *a subsequence of a frequent sequence is not necessarily frequent*. Several specialized approaches have been proposed [6,10,18] but they have a lack of genericity to handle simultaneously various types of constraints. Recently, a few proposals [4,8,11,12] have investigated relationships between GSPM and constraint programming (CP) in order to provide a declarative approach, while exploiting efficient and generic solving methods. But, due to the size of the proposed encodings, these CP methods are not as efficient as specialized ones. More recently, we have proposed the global constraint PREFIX-PROJECTION for SPM which remedies to this drawback [7]. However, as this global constraint uses the projected databases principle, it cannot be directly extended to support gap constraint.

In this paper, we introduce the global constraint GAP-SEQ enabling to handle SPM with or without gap constraint. GAP-SEQ relies on the principle of right pattern extension and its filtering exploits the prefix anti-monotonicity property of the gap constraint to provide an efficient pruning of the search space. GAP-SEQ enables to handle simultaneously different types of constraints and its encoding does not require any reified constraints nor any extra variables. Finally, experiments show that our approach clearly outperforms CP approaches as well as specialized methods for GSPM and achieves scalability while it is a major issue for CP approaches.

The paper is organized as follows. Section 2 introduces the prefix anti-monotonicity of the gap constraint as well as right pattern extensions that will enable an efficient filtering. Section 3 provides a critical review of specialized methods and CP approaches for sequential pattern mining under gap constraint. Section 4 presents the global constraint GAP-SEQ. Section 5 reports experiments we performed. Finally, we conclude and draw some perspectives.

2 Preliminaries

First, we provide the basic definitions for GSPM. Then, we show that the anti-monotonicity property of frequency of SPM does not hold for GSPM. Finally, we introduce right pattern extensions that will enable an efficient filtering for GSPM.

2.1 Definitions

Let $\mathcal{I}$ be a finite set of distinct *items*. The language of sequences corresponds to $\mathcal{L}_{\mathcal{I}} = \mathcal{I}^n$ where $n \in \mathbb{N}^+$.

Definition 1 (sequence, sequence database). *A sequence s over $\mathcal{L}_{\mathcal{I}}$ is an ordered list $\langle s_1 s_2 \dots s_n \rangle$, where s_i, $1 \leq i \leq n$, is an item. n is called the length of the sequence s. A sequence database SDB is a set of tuples (sid, s), where sid is a sequence identifier and s a sequence denoted by $SDB[sid]$.*

We now define the subsequence relation $\preceq^{[M,N]}$ under $gap[M,N]$ constraint which restricts the allowed distance between items of subsequences in sequences.

Definition 2 (subsequence relation $\preceq^{[M,N]}$ under $gap[M,N]$). *$\alpha = \langle \alpha_1 \dots \alpha_m \rangle$ is a subsequence of $s = \langle s_1 \dots s_n \rangle$, under $gap[M,N]$, denoted by $(\alpha \preceq^{[M,N]} s)$, if $m \leq n$ and, for all $1 \leq i \leq m$, there exist integers $1 \leq j_1 \leq \dots \leq j_m \leq n$, such that $\alpha_i = s_{j_i}$, and $\forall k \in \{1, ..., m-1\}, M \leq j_{k+1} - j_k - 1 \leq N$. In this context, the pair $(s, [j_1, j_m])$ denotes an* **occurrence** *of α in s, where j_1 and j_m represent the positions of the first and last items of α in s. We say that α is contained in s or s is a super-sequence of α under $gap[M,N]$. We also say that α is a $gap[M,N]$* **constrained pattern** *in s.*

- *Let $AllOcc(\alpha, s) = \{[j_1, j_m] \mid (s, [j_1, j_m])$ is an occurrence of α in $s\}$ be the set of all the occurrences of some sequence α under $gap[M,N]$ in s.*
- *Let $AllOcc(\alpha, SDB) = \{(sid, AllOcc(\alpha, SDB[sid])) \mid (sid, SDB[sid]) \in SDB\}$ be the set of all the occurrences of some sequence α under $gap[M,N]$ in SDB.*
- *Let $gap[M, \infty]$ and $gap[0, N]$ the* **minimum and the maximum gap** *constraints respectively.*
- *The relation $\preceq$ stands for $\preceq^{[0,\infty]}$ where the gap constraint is inactive.*

sid	Sequence
1	$\langle ABCDB \rangle$
2	$\langle ACCBACB \rangle$
3	$\langle ADCBEEC \rangle$
4	$\langle AACC \rangle$

Table 1: A sequence database example SDB_1.

For example, the sequence $\langle BABC \rangle$ is a super-sequence of $\langle AC \rangle$ under $gap[0,2]$: $\langle AC \rangle \preceq^{[0,2]} \langle BABC \rangle$.

Definition 3 (prefix, postfix). *Let $\beta = \langle \beta_1 \dots \beta_n \rangle$ be a sequence. The sequence $\alpha = \langle \alpha_1 \dots \alpha_m \rangle$ where $m \leq n$ is called the prefix of β iff $\forall i \in [1..m], \alpha_i = \beta_i$. The sequence $\gamma = \langle \beta_{m+1} \dots \beta_n \rangle$ is called the postfix of s w.r.t. α. With the standard concatenation operator "concat", we have $\beta = concat(\alpha, \gamma)$.*

The cover of a sequence α in SDB is the set of all tuples in SDB in which α is contained. The support of a sequence α in SDB is the cardinal of its cover.

Definition 4 (coverage and support under $gap[M,N]$). *Let SDB be a sequence database and α a sequence. $cover_{SDB}^{[M,N]}(\alpha) = \{(sid, s) \in SDB \mid \alpha \preceq^{[M,N]} s\}$ and $sup_{SDB}^{[M,N]}(\alpha) = \#cover_{SDB}^{[M,N]}(\alpha)$.*

Definition 5 ($gap[M,N]$ constrained sequential pattern mining (GSPM)). *Given a sequence database SDB, a minimum support threshold $minsup$ and a gap constraint $gap[M,N]$. The problem of $gap[M,N]$ constrained sequential pattern mining is to find all subsequences α such that $sup_{SDB}^{[M,N]}(\alpha) \geq minsup$.*

Example 1. Table 1 represents a sequence database of four sequences where the set of items is $\mathcal{I} = \{A, B, C, D, E\}$. Let the sequence $\alpha = \langle AC \rangle$. The occurrences under $gap[0,1]$ of α in $SDB_1[2]$ is given by $AllOcc(\alpha, SDB_1[2]) = \{[1,2]), [1,3], [5,6]\}$. We have $cover_{SDB_1}^{[0,1]}(\alpha) = \{(1, s_1), (2, s_2), (3, s_3), (4, s_4)\}$. If we consider $minsup = 2$, α is a $gap[0,1]$ constrained sequential pattern because $sup_{SDB_1}^{[0,1]}(\alpha) \geq 2$.

2.2 Prefix anti-monotonicity of $gap[M, N]$

Most of SPM algorithms rely on the *anti-monotonicity property of frequency* [1] to reduce the search space: all the subsequences of a frequent sequence are frequent as well (or, equivalently, if a subsequence is infrequent, then no super-sequence of it can be frequent). However, this property does not hold for the gap constraint, and more precisely for the maximum gap constraint. A simple illustration from our running example suffices to show that sequence $\langle AB \rangle$ is not a sequential pattern under $gap[0,1]$ (for $minsup = 3$) whereas sequence $\langle ACB \rangle$ is a $gap[0,1]$ constrained sequential pattern. As a consequence, one needs to use other techniques for pruning the search space. The following proposition shows how the *prefix anti-monotonicity property* can be exploited to ensure the anti-monotonicity of the gap constraint.

Definition 6 (prefix anti-monotone property). *A constraint c is called prefix anti-monotone if for every sequence α satisfying c, every prefix of α also satisfies the constraint.*

Proposition 1. *$gap[M, N]$ is prefix anti-monotone.*

Proof. Let $\alpha = \langle \alpha_1 \dots \alpha_m \rangle$ and $s = \langle s_1 \dots s_n \rangle$ be two sequences s.t. $\alpha \preceq^{[M,N]} s$ and $m \leq n$. By definition, there exist integers $1 \leq j_1 \leq \dots \leq j_m \leq n$, such that $\alpha_i = s_{j_i}$, and $\forall k \in \{1, ..., m-1\}, M \leq j_{k+1} - j_k - 1 \leq N$. As a consequence, the property also holds for every prefix of α. □

Hence, if a sequence α does not satisfy $gap[M, N]$, then all sequences that have α as prefix will not satisfy this constraint. Sect. 4.2 shows how this property can be exploited to provide an efficient filtering.

2.3 Right pattern extensions

Right pattern extensions of some pattern p gives all the possible subsequences which can be appended at right of p to form a $gap[M, N]$ constrained pattern. According to proposition 1, the set of all items locally frequent within the right pattern extensions of p in SDB can be used to extend p. In the following, we introduce an operator allowing to compute all the right pattern extensions of a pattern w.r.t. $gap[M, N]$.

Definition 7 (Right pattern extensions). *Given some sequence (sid, s) and a pattern p s.t. $p \preceq^{[M,N]} s$. The* **right pattern extensions** *of p in s, denoted by $Ext_R^{[M,N]}(p, s)$, is the collection of legal subsequences of s located at the right of p and satisfying $gap[M, N]$. To define $Ext_R^{[M,N]}(p, s)$, we need to define $BE^{[M,N]}(p, s)$* **basic right extensions** *:*

$$BE^{[M,N]}(p,s) = \bigcup_{[j_1,j_m] \in AllOcc(p,s)} \{(j_m, \texttt{SubSeq}(s, j_m + M + 1, min(j_m + N + 1, \#s)))\}$$

$$\textit{where } \texttt{SubSeq}(s, i_1, i_2) = \begin{cases} \langle s[i_1], ..., s[i_2] \rangle & \textit{if } i_1 \leq i_2 \leq \#s \\ \langle\rangle & \textit{otherwise} \end{cases}$$

Right pattern extensions $Ext_R^{[M,N]}(p, s)$ is defined as follows:

$$Ext_R^{[M,N]}(p,s) = \begin{cases} \{Sb \mid (j'_m, Sb) \in BE^{[M,N]}(p,s) \wedge & \textit{if } N \geq \#s \\ \quad j'_m = \min_{(j_m,Sb) \in BE^{[M,N]}(p,s)} \{j_m\}\} & \\ \bigcup_{(j_m,Sb) \in BE^{[M,N]}(p,s)} \{Sb\} & \textit{otherwise} \end{cases} \quad (1)$$

Formula (1) states exactly the set of all possible extensions of pattern p within s. In case where $(N \geq \#s)$, since that any extension from $BE^{[M,N]}(p, s)$ always reaches the end of the sequence s, thus all possible extensions can be aggregated within one unique extension going from the lowest starting position $j'_m = \min_{(j_m,Sb) \in BE^{[M,N]}(p,s)} \{j_m\}$. We point out that these cases $(N \geq \#s)$ cover the special case of no gap $gap[0, \infty]$.

The right pattern extensions of p in SDB is the collection of all its right pattern extensions in all sequences of SDB:

$$Ext_R^{[M,N]}(p, SDB) = \bigcup_{(sid,s) \in SDB} \{(sid, Ext_R^{[M,N]}(p, s))\} \quad (2)$$

Example 2. Let $p_1 = \langle AC \rangle$ be a pattern and the gap constraint be $gap[0,1]$. We have $AllOcc(p_1, SDB_1[2]) = \{[1,2]), [1,3], [5,6]\}$. The right pattern extensions of p_1 in $SDB_1[2]$ is equal to $Ext_R^{[0,1]}(p_1, SDB_1[2]) = \{\langle CB\rangle, \langle BA\rangle, \langle B\rangle\}$. The right pattern extensions of p_1 in SDB_1 is given by $Ext_R^{[0,1]}(p_1, SDB_1) = \{(1, \{\langle DB\rangle\}), (2, \{\langle CB\rangle, \langle BA\rangle, \langle B\rangle\}), (3, \{\langle BE\rangle\}), (4, \{\langle C\rangle\})\}$.

Let the gap constraint be $gap[0,\infty]$. To compute $Ext_R^{[0,\infty]}(p_1, SDB_1[2])$, only the first occurrence of p_1 in $SDB_1[2]$ need to be considered (i.e. $[1,2]$) (cf. Definition 7). Thus, $Ext_R^{[0,\infty]}(p_1, SDB_1[2]) = \{\langle CBACB\rangle\})$. The right pattern extensions of p_1 in SDB_1 is equal to $Ext_R^{[0,\infty]}(p_1, SDB_1) = \{(1, \{\langle DB\rangle\}), (2, \{\langle CBACB\rangle\}), (3, \{\langle BEEC\rangle\}), (4, \{\langle C\rangle\})\}$.

Given a $gap[M,N]$ constrained pattern p in SDB, according to proposition 1, the set of all items locally frequent within the right pattern extensions of p in SDB can be used to extend p. Proposition 2 establishes the support count of a sequence γ w.r.t. its right pattern extensions.

Proposition 2 (Support count). *For any sequence γ in SDB with prefix α and postfix β s.t. $\gamma = concat(\alpha, \beta)$, $sup_{SDB}^{[M,N]}(\gamma) = sup_{Ext_R^{[M,N]}(\alpha, SDB)}(\beta)$.*

This proposition ensures that only the sequences in SDB grown from α need to be considered for the support count of a sequence γ. From proposition 2, we can derive the following proposition to establish a condition to check when a pattern is a $gap[M,N]$ constrained sequential pattern.

Proposition 3. *Let SDB be a sequence database and a minimum support threshold $minsup$. A pattern p is a $gap[M,N]$ constrained sequential pattern in SDB if and only if the following condition holds: $\#Ext_R^{[M,N]}(p, SDB) \geq minsup$*

Example 3. Let $minsup$ be 2 and the gap constraint be $gap[0,1]$. From Example 2, we have $\#Ext_R^{[0,1]}(p_1, SDB_1) = 4 \geq minsup$. Thus, $p_1 = \langle AC\rangle$ is a $gap[0,1]$ constrained sequential pattern. The locally frequent items within the right pattern extensions $Ext_R^{[0,1]}(p_1, SDB_1)$ of p_1 are B and C with supports of 3 and 2 respectively. According to proposition 2, p_1 can be extended to two $gap[0,1]$ constrained sequential patterns $\langle ACB\rangle$ and $\langle ACC\rangle$.

3 Related works

Specialized methods for GSPM. The SPM was first proposed in [1]. Since then, many efficient specialized approaches have been proposed [2,13,17]. There are also several methods focusing on gap constraints. Zaki [16] first proposed `cSpade`, a depth-first search based on a vertical database format, incorporating max-gap, max-span and length constraints. Ji and al. [6] and Li and al. [9] studied the problem of mining frequent patterns with gap constraints. In [6], a minimal distinguishing subsequence that occurs frequently in the positive sequences and infrequently in the negative sequences is proposed, where the maximum gap constraint is defined. In [9], closed frequent patterns with gap constraints are mined. All these proposals, though efficient, lack of genericity to handle simultaneously various types of constraints.

CP Methods for GSPM. There are few methods for SPM with gap constraints using CP. [11] have proposed to model a sequence using an automaton capturing all subsequences that can occur in it. The gap constraint is encoded by removing from the automaton all transitions that does not respect the gap constraint. [8] have proposed a CSP model for SPM with explicit wildcards[1]. The gap constraints is enforced using the regular global constraint. [12] have proposed two CP encodings for the SPM. The first one uses a global constraint to encode the subsequence relation (denoted `global-p.f`), while the second one (denoted `decomposed-p.f`) encodes explicitly this relation using additional variables and constraints in order to support constraints like gap. However, all these proposals usually lead to constraints network of huge size. Space complexity is clearly identified as the main bottleneck behind the competitiveness of these declarative approaches. In [7], we have proposed the global constraint PREFIX-PROJECTION for sequential pattern mining which remedies to this drawback. However, this constraint cannot be directly extended to handle gap constraints. This requires changing the way the subsequence relation is encoded.

The next section introduces the global constraint GAP-SEQ enabling to handle SPM with or without gap constraints. GAP-SEQ relies on the prefix anti-monotonicity of the gap constraint and on the right pattern extensions to provide an efficient filtering. This global constraint does not require any reified constraints nor any extra variables to encode the subsequence relation.

4 GAP-SEQ global constraint

This section is devoted to the GAP-SEQ global constraint. Section 4.1 defines the GAP-SEQ global constraint and presents the CSP modeling. Section 4.2 shows how the filtering can take advantage of the prefix anti-monotonicity property of the $gap[M, N]$ constraint (see Proposition 6) and of the right pattern extensions (see Proposition 5) to remove inconsistent values from the domain of a future variable. Section 4.3 details the filtering algorithm and Section 4.4 provides its temporal and spatial complexities.

4.1 CSP modeling for GSPM

A *Constraint Satisfaction Problem* (CSP) consists of a set X of n variables, a domain $\mathcal{D}$ mapping each variable $X_i \in X$ to a finite set of values $D(X_i)$, and a set of constraints $\mathcal{C}$. An assignment σ is a mapping from variables in X to values in their domains. A constraint $c \in \mathcal{C}$ is a subset of the cartesian product of the domains of the variables that occur in c. The goal is to find an assignment such that all constraints are satisfied.

(a) Variables and domains. Let P be the unknown pattern of size ℓ we are looking for. The symbol $\square$ stands for an empty item and denotes the end of a sequence. We encode the unknown pattern P of maximum length ℓ with a sequence of ℓ variables $\langle P_1, P_2, \ldots, P_\ell \rangle$. Each variable P_j represents the item in the jth position of the sequence. The size ℓ of P is determined by the length of the longest sequence of SDB. The domains of variables are defined as follows: (i) $D(P_1) = \mathcal{I}$ to avoid the empty sequence, and (ii) $\forall i \in [2 \ldots \ell], D(P_i) = \mathcal{I} \cup \{\square\}$. To allow patterns with less than ℓ items, we impose that $\forall i \in [2..(\ell-1)], (P_i = \square) \rightarrow (P_{i+1} = \square)$.

[1] A wildcard is a special symbol that matches any item of $\mathcal{I}$ including itself.

(b) Definition of GAP-SEQ. The global constraint GAP-SEQ encodes both subsequence relation $\preceq^{[M,N]}$ under gap constraint $gap[M,N]$ and minimum frequency constraint directly on the data.

Definition 8 (GAP-SEQ global constraint). *Let $P = \langle P_1, P_2, \ldots, P_\ell \rangle$ be a pattern of size ℓ and $gap[M,N]$ be the gap constraint. $\langle d_1, ..., d_\ell \rangle \in D(P_1) \times \ldots \times D(P_\ell)$ is a solution of* GAP-SEQ$(P, SDB, minsup, M, N)$ *iff $sup_{SDB}^{[M,N]}(\langle d_1, ..., d_\ell \rangle) \geq minsup$.*

Proposition 4. GAP-SEQ$(P, SDB, minsup, M, N)$ *has a solution iff there exists an assignment $\sigma = \langle d_1, ..., d_\ell \rangle$ of variables of P s.t. $\#Ext_R^{[M,N]}(\sigma, SDB) \geq minsup$.*

Proof: This is a direct consequence of proposition 3. □

(c) Other SPM constraints can be directly modeled as follows:
- *Minimum Size* constraint restricts the number of items of a pattern to be at least ℓ_{min}: $minSize(P, \ell_{min}) \equiv \bigwedge_{i=1}^{i=\ell_{min}} (P_i \neq \square)$
- *Maximum Size* constraint restricts the number of items of a pattern to be at most ℓ_{max}: $maxSize(P, \ell_{max}) \equiv \bigwedge_{i=\ell_{max}+1}^{i=\ell} (P_i = \square)$
- *Membership* constraint states that a subset of items V must belong (or not) to the extracted patterns. $item(P, V) \equiv \bigwedge_{t \in V} \texttt{Among}(P, \{t\}, l, u)$ enforces that items of V should occur at least l times and at most u times in P. To forbid items of V to occur in P, l and u must be set to 0.

4.2 Principles of filtering

(a) Maintaining a local consistency. SPM is a challenging task due to the exponential number of candidates that should be parsed to find the frequent patterns. For instance, with k items there are $O(n^k)$ potentially candidate patterns of length at most k in a sequence of size n. With gap constraints, the problem is even much harder since the complexity of checking for subsequences taking a gap constraint into account is higher than the complexity of the standard subsequence relation. Furthermore, the NP-hardness of mining maximal[2] frequent sequences was established in [15] by proving the #P-completeness of the problem of counting the number of maximal frequent sequences. Hence, ensuring *Domain Consistency* (DC) for GAP-SEQ i.e., finding, for every variable P_j, a value $d_j \in D(P_j)$, satisfying the constraint is NP-hard.

So, the filtering of GAP-SEQ constraint maintains a consistency lower than DC. This consistency is based on specific properties of the $gap[M,N]$ constraint and resembles forward-checking (regarding Proposition 5). GAP-SEQ is considered as a global constraint, since all variables share the same internal data structures that awake and drive the filtering. The prefix anti-monotonicity property of the $gap[M,N]$ constraint (see Proposition 6) and of the right pattern extensions (see Proposition 5) will enable to remove inconsistent values from the domain of a future variable.

(b) Detecting inconsistent values. Let $\mathcal{RF}^{[M,N]}(\sigma, SDB)$ be the set of locally frequent items within the right pattern extensions, defined by $\{v \in \mathcal{I} \,|\, \#\{sid \,|\, (sid, E) \in Ext_R^{[M,N]}(\sigma, SDB) \wedge (\exists \alpha \in E \wedge \langle v \rangle \preceq \alpha)\} \geq minsup\}$. The following proposition

[2] A sequential pattern p is maximal if there is no sequential pattern q such that $p \preceq q$.

characterizes values, of a future (unassigned) variable P_{j+1}, that are consistent with the current assignment of variables $\langle P_1, \ldots, P_j \rangle$.

Proposition 5. *Let*[3] $\sigma = \langle d_1, \ldots, d_j \rangle$ *be a current assignment of variables* $\langle P_1, \ldots, P_j \rangle$, P_{j+1} *be a future variable. A value* $d \in D(P_{j+1})$ *occurs in a solution for the global constraint* GAP-SEQ$(P, SDB, minsup, M, N)$ *iff* $d \in \mathcal{RF}^{[M,N]}(\sigma, SDB)$.

Proof: Assume that $\sigma = \langle d_1, \ldots, d_j \rangle$ is $gap[M, N]$ constrained sequential pattern in SDB. Suppose that value $d \in D(P_{j+1})$ appears in $\mathcal{RF}^{[M,N]}(\sigma, SDB)$. As the local support of d within the right extensions is equal to $sup_{Ext_R^{[M,N]}(\sigma,SDB)}(\langle d \rangle)$, from proposition 2 we have $sup_{SDB}^{[M,N]}(concat(\sigma, \langle d \rangle)) = sup_{Ext_R^{[M,N]}(\sigma,SDB)}(\langle d \rangle)$. Hence, we can get a new assignment $\sigma \cup \langle d \rangle$ that satisfies the constraint. Therefore, $d \in D(P_{j+1})$ participates in a solution. □

From proposition 5 and according to the prefix anti-monotonicity property of the gap constraint, we can derive the following pruning rule:

Proposition 6. *Let* $\sigma = \langle d_1, \ldots, d_j \rangle$ *be a current assignment of variables* $\langle P_1, \ldots, P_j \rangle$. *All values* $d \in D(P_{j+1})$ *that are not in* $\mathcal{RF}^{[M,N]}(\sigma, SDB)$ *can be removed from the domain of variable* P_{j+1}.

Example 4. Consider the running example of Table 1, let $minsup$ be 2 and the gap constraint be $gap[1, 2]$. Let $P = \langle P_1, P_2, P_3, P_4 \rangle$ with $D(P_1) = \mathcal{I}$ and $D(P_2) = D(P_3) = D(P_4) = \mathcal{I} \cup \{\square\}$. Suppose that $\sigma(P_1) = A$. We have $Ext_R^{[1,2]}(\langle A \rangle, SDB_1) = \{(1, \{\langle CD \rangle\}), (2, \{\langle CB \rangle, \langle B \rangle\}), (3, \{\langle CB \rangle\}), (4, \{\langle CC \rangle, \langle C \rangle\})\}$. As B and C are the only locally frequent items in $Ext_R^{[1,2]}(\langle A \rangle, SDB_1)$, GAP-SEQ will remove values A, D and E from $D(P_2)$.

The filtering of GAP-SEQ relies on Proposition 6 and is detailed in the next section.

4.3 Filtering algorithm

Algorithm 1 describes the pseudo-code of GAP-SEQ filtering algorithm. It takes as input: the index j of the last assigned variable in P, the current partial assignment $\sigma = \langle \sigma(P_1), \ldots, \sigma(P_j) \rangle$, the minimum support threshold $minsup$, the minimum and the maximum gaps. The internal data-structure $\mathcal{ALLOCC}$ stores all the intermediate occurrences of patterns in SDB, where $\mathcal{ALLOCC}_j = AllOcc(\sigma, SDB)$, for $j \in [1 \ldots \ell]$. If $\sigma = \langle\rangle$, then $\mathcal{ALLOCC}_0 = \{(sid, [1, \#s]) \mid (sid, s) \in SDB\}$. $\mathcal{ALLOCC}_j$ is computed incrementally from $\mathcal{ALLOCC}_{j-1}$ in order to enhance the efficiency.

Algorithm 1 starts by computing the right pattern extensions Ext_R of σ in SDB (line 1) by calling function GETRIGHTEXT (see Algorithm 2). Then, it checks whether the current assignment σ satisfies the constraint (see Proposition 4) (line 2). If not, we stop growing σ and return *False*. Otherwise, the algorithm checks if the last assigned variable P_j is instantiated to $\square$ (line 4). If so, the end of the sequence is reached (since value $\square$ can only appear at the end) and the sequence $\langle \sigma(P_1), \ldots, \sigma(P_j) \rangle$ is a

[3] We indifferently denote σ by $\langle d_1, \ldots, d_j \rangle$ or by $\langle \sigma(P_1), \ldots, \sigma(P_j) \rangle$.

Algorithm 1: FILTER-GAP-SEQ(SDB, σ, j, P, $minsup$, M, N)

Data: SDB: initial database; σ: current assignment $\langle \sigma(P_1), \ldots, \sigma(P_j) \rangle$; $minsup$: the minimum support threshold; $\mathcal{ALLOCC}$: internal data structure for storing occurrences of patterns in SDB; Ext_R: internal data structure for storing right pattern extensions of σ in SDB.

begin

1 $Ext_R \leftarrow$ GETRIGHTEXT($SDB, \mathcal{ALLOCC}_{j-1}, \sigma, M, N$) ;
2 **if** ($\#Ext_R < minsup$) **then**
3 **return** False ;
4 **if** ($j \geq 2 \wedge \sigma(P_j) = \square$) **then**
5 **for** $k \leftarrow j+1$ **to** ℓ **do**
6 $P_k \leftarrow \square$;
else
7 $\mathcal{RF} \leftarrow$ GETFREQITEMS($SDB, Ext_R, minsup$) ;
8 **foreach** $a \in D(P_{j+1})\ s.t. (a \neq \square \wedge a \notin \mathcal{RF})$ **do**
9 $D(P_{j+1}) \leftarrow D(P_{j+1}) - \{a\}$;
10 **return** True ;

$gap[M, N]$ constrained sequential pattern in SDB; hence, the algorithm sets the remaining $(\ell - j)$ unassigned variables to $\square$ and returns *True* (lines 5-6). If $(P_j \neq \square)$, the set of locally frequent items, within the right pattern extensions Ext_R of σ in SDB, is computed by calling function GETFREQITEMS (line 7) and the domain of variable P_{j+1} is updated accordingly (lines 8-9).

Algorithm 2 gives the pseudo-code of the function GETRIGHTEXT. First, if σ is empty (i.e. $\#\sigma = 0$), all the sequences of SDB are considered as valid right pattern extensions; the whole SDB should be returned. Otherwise, the function GETALLOCC is called to compute the occurrences of σ in SDB (line 3). Then, the algorithm processes all the entries of $\mathcal{ALLOCC}_j$, one by one (line 5), and, for each pair $(sid, OccSet)$, scans the occurrences of σ in the sequence sid (line 7). For each occurrence $(j_1, j_m) \in OccSet$, the algorithm computes its right pattern extensions, i.e. the part of the sequence sid which is in the range $[j_m + M + 1, min(j_m + N + 1, \#s)]$ (line 8). If the new range is valid, it is added to the set Sb (line 10). After processing the whole entries in $OccSet$, the right pattern extensions of σ in the sequence sid are built and then added to the set Ext_R (line 11). The process ends when all entries of $\mathcal{ALLOCC}_j$ have been considered. The right pattern extensions of σ in SDB are then returned (line 12).

The function GETALLOCC computes incrementally $\mathcal{ALLOCC}_j$ from $\mathcal{ALLOCC}_{j-1}$. More precisely, lines (18-19) and (24-25) are considered when the first variable P_1 is instanciated (i.e. $\#\sigma = 1$), and consequently all of its initial occurrences should be found and stored in $\mathcal{ALLOCC}_1$ through the initialization step (lines 24-25). After that, $\mathcal{ALLOCC}_j (j > 1)$ is incrementally computed from $\mathcal{ALLOCC}_{j-1}$ through line (26). To determine $\mathcal{ALLOCC}_j$, we avoid computing occurrences leading to redundant right pattern extensions thanks to the conditions $((sup = \#s) \wedge (\#\sigma > 1))$ in line (27). Moreover, when computing the right pattern extensions, instead of storing the part of subsequence $\langle s[j'_1], \ldots, s[j'_m] \rangle$, one can only store the positions of its first and last items (j'_1, j'_m) in the sequence sid. Consider Example 2: $Ext_R^{[0,1]}(\langle AC \rangle, SDB_1)$ will be encoded as $\{(1, \{(4,5)\}), (2, \{(3,4), (4,5), (7,7)\}), (3, \{(4,5)\}), (4, \{(4,4)\})\}$.

Finally, the filtering algorithm handles as efficiently the case *without gap constraints*. For each pair $(sid, OccSet)$, only the first occurrence (j_1, j_m) in $OccSet$ is determined thanks to the condition $(N \geq \#s)$ in line (27).

Algorithm 2: GETRIGHTEXT(SDB, $\mathcal{ALLOCC}_{j-1}$, σ, M, N)

Data: SDB: initial database; $\mathcal{ALLOCC}_{j-1}$: occurrences of the partial assignment $\langle\sigma(P_1), \ldots, \sigma(P_{j-1})\rangle$ in SDB; σ: the current partial assignment $\langle\sigma(P_1), \ldots, \sigma(P_j)\rangle$; $OccSet$: the positions of the first and last items of $\langle\sigma(P_1), \ldots, \sigma(P_{j-1})\rangle$ in $SDB[sid]$; Sb: the positions of the first and last items of the right pattern extensions of σ in $SDB[sid]$.

begin

if $(\sigma = \langle\rangle)$ **then**
  return $\{(sid, [1, \#s]) | (sid, s) \in SDB\}$;
$\mathcal{ALLOCC}_j \leftarrow$ GETALLOCC$(SDB, \mathcal{ALLOCC}_{j-1}, \sigma, M, N)$;
$Ext_R \leftarrow \emptyset$;
foreach *pair* $(sid, OccSet) \in \mathcal{ALLOCC}_j$ **do**
  $s \leftarrow SDB[sid]$; $Sb \leftarrow \emptyset$;
  foreach *pair* $[j_1, j_m] \in OccSet$ **do**
    $j'_1 \leftarrow j_m + M + 1$; $j'_m \leftarrow min(j_m + N + 1, \#s)$;
    if $(j'_1 \leq j'_m)$ **then**
      $Sb \leftarrow Sb \cup \{(j'_1, j'_m)\}$;
  $Ext_R \leftarrow Ext_R \cup \{(sid, Sb)\}$;
return Ext_R ;

FUNCTION GETALLOCC (SDB, $\mathcal{ALLOCC}_{j-1}$, σ, M, N) ;

begin

$\mathcal{ALLOCC}_j \leftarrow \emptyset$; $inf \leftarrow 0$; $sup \leftarrow 0$;
foreach *pair* $(sid, OccSet) \in \mathcal{ALLOCC}_{j-1}$ **do**
  $s \leftarrow SDB[sid]$; $newOccSet \leftarrow \emptyset$; $redundant \leftarrow false$; $i \leftarrow 1$;
  while $(i \leq \#OccSet \wedge \neg redundant)$ **do**
    $[j_1, j_m] \leftarrow OccSet[i]$; $i \leftarrow i + 1$;
    if $(\#\sigma = 1)$ **then**
      $inf \leftarrow 1$; $sup \leftarrow \#s$;
    else
      $inf \leftarrow j_m + M + 1$; $sup \leftarrow min(j_m + N + 1, \#s)$;
    $k \leftarrow inf$;
    while $((k \leq sup) \wedge (\neg redundant))$ **do**
      if $(s[k] = \sigma(P_j))$ **then**
        if $(\#\sigma = 1)$ **then**
          $newOccSet \leftarrow newOccSet \cup \{[k, k]\}$;
        else
          $newOccSet \leftarrow newOccSet \cup \{[j_1, k]\}$;
        if $(((sup = \#s) \wedge (\#\sigma > 1)) \vee (N \geq \#s))$ **then**
          $redundant \leftarrow true$;
      $k \leftarrow k + 1$;
  if $(newOccSet \neq \emptyset)$ **then**
    $\mathcal{ALLOCC}_j \leftarrow \mathcal{ALLOCC}_j \cup (sid, newOccSet)$;
return $\mathcal{ALLOCC}_j$;

4.4 Temporal and spatial complexities of the filtering algorithm

Let m=$|SDB|$, d=$|\mathcal{I}|$, and ℓ be the length of the longest sequence in SDB. Computing $\mathcal{ALLOCC}_j$ from $\mathcal{ALLOCC}_{j-1}$ (see function GETALLOCC of Algorithm 2) can be achieved in $O(m \times \ell^2)$. The function GETRIGHTEXT (see Algorithm 2) processes all the occurrences of σ in each sequence of the SDB. In the worst case, it may exist ℓ occurrences for each sequence in the database. So, the time complexity of function GETRIGHTEXT is $O(m \times \ell^2 + m \times \ell)$ i.e. $O(m \times \ell^2)$.

Proposition 7. *In the worst case, (i) filtering can be achieved in $O(m \times \ell^2 + d)$ and (ii) the space complexity is $O(m \times \ell^2)$.*

dataset	$\#SDB$	$\#\mathcal{I}$	avg $(\#s)$	$\max_{s \in SDB}(\#s)$	type of data
Leviathan	5834	9025	33.81	100	book
PubMed	17527	19931	29	198	bio-medical text
FIFA	20450	2990	34.74	100	web click stream
BIBLE	36369	13905	21.64	100	bible
Kosarak	69999	21144	7.97	796	web click stream
Protein	103120	24	482	600	protein sequences

Table 2: Dataset Characteristics.

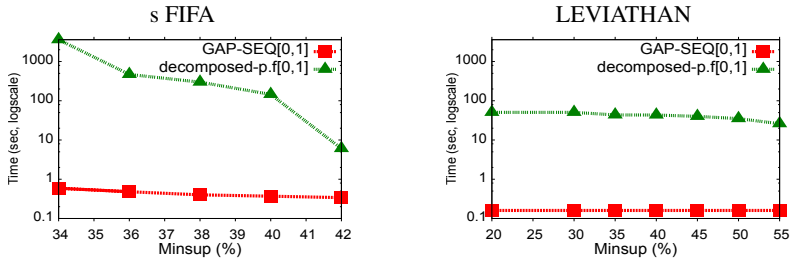

Fig. 1: Comparing `GAP-SEQ` with `decomposed-p.f` for GSPM: CPU times.

Proof: (i) The complexity of function GETRIGHTEXT is $O(m \times \ell^2)$. The total complexity of function GETFREQITEMS is $O(m \times \ell)$. Lines (8-9) can be achieved in $O(d)$. So, the whole complexity is $O(m \times \ell^2 + m \times \ell + d)$, i.e. $O(m \times \ell^2 + d)$.
(ii) The space complexity of the filtering algorithm lies in the storage of the $\mathcal{ALLOCC}$ internal data structure. The occurrences $\mathcal{ALLOCC}_j$ of each assignment σ in SDB, with the length of σ varying from 1 to ℓ, have to be stored. Since it may exist at most ℓ occurrences of σ in each sequence sid, storing any $\mathcal{ALLOCC}_j$ costs in the worst case $O(m \times \ell)$. Since we can have ℓ prefixes, the worst space complexity of storing all the occurrences $\mathcal{ALLOCC}_j (j = 1..\ell)$, is $O(m \times \ell^2)$. □

5 Experiments

This section reports experiments on several real-life datasets [5,3] of large size having varied characteristics and representing different application domains (see Tab. 2). First, we compare our approach with CP methods and with the state-of-the-art specialized method `cSpade` in terms of scalability. Second, we show the flexibility of our approach for handling different types of constraints simultaneously.

Experimental protocol. Our approach was carried out using the `gecode` solver[4]. All experiments were conducted on a processor Intel X5670 with 24 GB of memory. A time limit of 1 hour has been set. If an approach is not able to complete the extraction within the time limit, it will be reported as $(-)$. ℓ was set to the length of the longest sequence of SDB. We compare our approach (indicated by `GAP-SEQ`) with:

1. `decomposed-p.f`[5], the most efficient CP methods for GSPM,
2. `cSpade`[6], the state-of-the-art specialized method for GSPM,

[4] `http://www.gecode.org`
[5] `https://dtai.cs.kuleuven.be/CP4IM/cpsm/`
[6] `http://www.cs.rpi.edu/~zaki/www-new/pmwiki.php/Software/`

Dataset	$minsup$ (%)	#PATTERNS	CPU times (s)		#PROPAGATIONS		#NODES	
			GAP-SEQ	decomposed-p.f	GAP-SEQ	decomposed-p.f	GAP-SEQ	decomposed-p.f
FIFA	42	1	0.34	6.06	2	0	1	2
	40	5	0.37	144.95	10	778010	6	11
	38	10	0.4	298.68	20	2957965	11	21
	36	17	0.48	469.3	34	9029578	18	35
	34	35	0.59	–	70	–	36	–

Table 3: GAP-SEQ vs. decomposed-p.f on FIFA dataset.

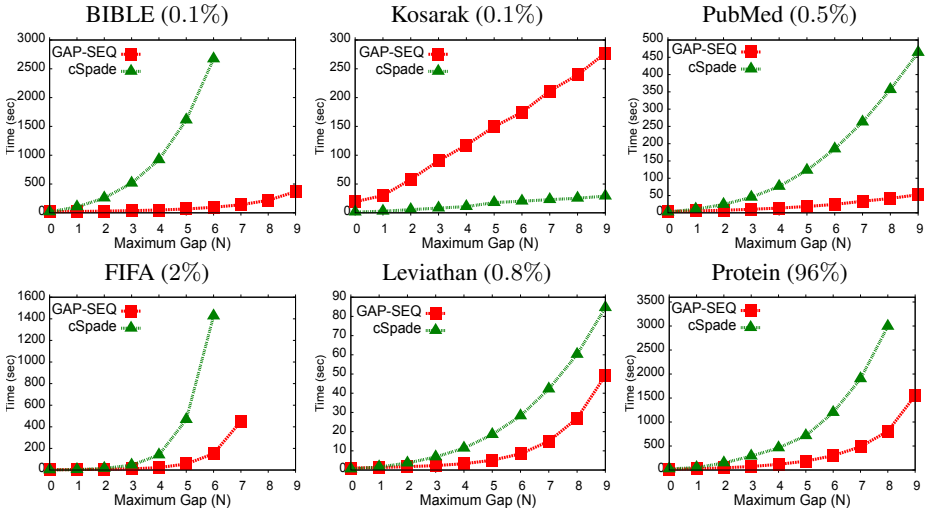

Fig. 2: Varying the value of parameter N in the gap constraint $(M = 0)$: CPU times.

3. the PREFIX-PROJECTION global constraint for SPM.

(a) GSPM: GAP-SEQ vs the most efficient CP method. We compare CPU times for GAP-SEQ and decomposed-p.f. In the experiments, we used the gap constraint $gap[0, 1]$ and various values of $minsup$. Fig. 1 shows the results for the two datasets FIFA and LEVIATHAN (results are similar for other datasets and not reported due to page limitation). GAP-SEQ clearly outperforms decomposed-p.f on the two datasets even for high values of $minsup$: GAP-SEQ is more than an order of magnitude faster than decomposed-p.f. For low values of $minsup$, decomposed-p.f fails to complete the extraction within the time limit.

Tab. 3 reports for the FIFA dataset and different values of $minsup$, the number of calls to the propagate function of gecode (col. 5) and the number of nodes of the search tree (col. 6). GAP-SEQ is very effective in terms of number of propagations. For GAP-SEQ, the number of propagations remains very small compared to decomposed-p.f (millions). This is due to the huge number of reified constraints used by decomposed-p.f to encode the subsequence relation. Regarding CPU times, GAP-SEQ requires less than 1s. to complete the extraction, while decomposed-p.f needs much more time to end the extraction (speed-up value up to 938).

(b) GSPM: GAP-SEQ vs the state-of-the-art specialized method. Second experiments compare GAP-SEQ with cSpade. We first fixed $minsup$ to the smallest possible value w.r.t. the dataset used, and varied the maximum gap N from 0 to 9. The

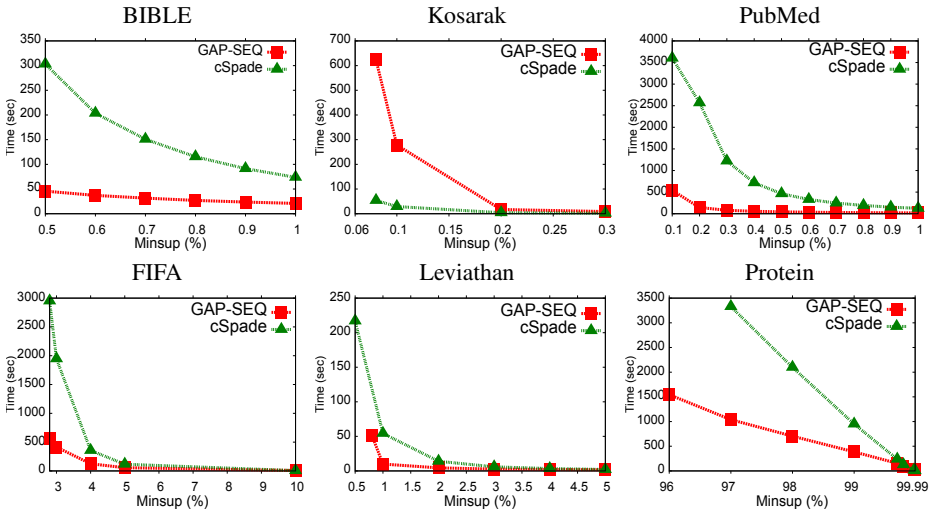

Fig. 3: Varying the value of $minsup$ with the gap constraint $gap[0, 9]$: CPU times.

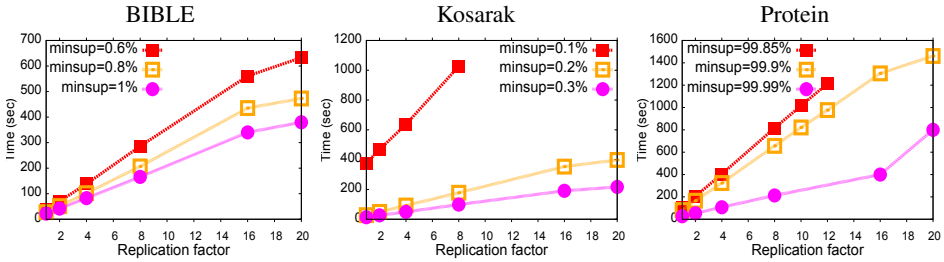

Fig. 4: Scalability of GAP−SEQ global constraint on BIBLE, Kosarak and Protein.

minimum gap M was set to 0. Fig. 2 reports the CPU times of both methods. First, GAP−SEQ clearly dominates cSpade on all the datasets. The gains in terms of CPU times are greatly amplified as the value of N increases. On FIFA, the speed-up is 9.5 for N=6. On BIBLE, GAP−SEQ is able to complete the extraction for values of N up to 9 in 433 seconds, while cSpade failed to complete the extraction for N greater than 6. The only exception is for the Kosarak dataset, where cSpade is efficient. For this dataset (which is the largest one both in terms of number of sequences and items), the size of the domains is important as compared to the other datasets. So, filtering takes much more time. This probably explains the behavior of GAP−SEQ on this dataset.

We also conducted experiments to evaluate how sensitive GAP−SEQ and cSpade are to $minsup$. We used the $gap[0, 9]$ constraint, while $minsup$ varied until the two methods were not able to complete the extraction within the time limit. Results are depicted in Fig. 3. Once again, GAP−SEQ obtains the best performance on all datasets (except for Kosarak). When the minimum support decreases, CPU times for GAP−SEQ increase reasonably while for cSpade they increase dramatically. On PubMed, with $minsup$ set to 0.1%, cSpade finished the extraction after $3,500$ seconds, while GAP−SEQ only used 500 seconds (speed-up value 7). These results clearly demonstrate that our approach is very effective as compared to cSpade on large datasets.

$minsup$	#PATTERNS		CPU times (s)		#PROPAGATIONS		#NODES	
	gap	gap+size+item	gap	gap+size+item	gap	gap+size+item	gap	gap+size+item
1 %	14032	1805	19.34	16.83	28862	47042	17580	16584
0.5 %	48990	6659	43.46	34.6	100736	163205	61149	58625
0.4 %	72228	10132	55.66	43.47	148597	240337	90477	87206
0.3 %	119965	17383	79.88	59.28	246934	398626	151280	146601
0.2 %	259760	39140	143.91	100.09	534816	861599	329185	321304
0.1 %	963053	153411	539.57	379.04	1986464	3186519	1236340	1219193

Table 4: GAP−SEQ under size and membership constraints on the PUBMED dataset.

(c) GSPM: evaluating the scalability of GAP−SEQ. We used three datasets and replicated them from 1 to 20 times. The gap constraint was set to $gap[0, 9]$, and $minsup$ to three different values. Fig. 4 reports the CPU times according to the replication factor (i.e. dataset sizes). CPU times increase (almost) linearly as the number of sequences. This indicates that GAP−SEQ achieves scalability while it is a major issue for CP approaches. The behavior of GAP−SEQ on Protein is quite different for low values of $minsup$. Indeed, for large sequences (such as in Protein), the size of $\mathcal{ALLOCC}$ may be very large and thus checking the gap constraint becomes costly (see Sect. 4.4).

(d) GSPM: handling various additional constraints. To illustrate the flexibility of our approach, we selected the PubMed dataset and stated additional constraints such as minimum frequency, minimum size, and other useful constraints expressing some linguistic knowledge as membership. The goal is to extract sequential patterns which convey linguistic regularities (e.g., gene - rare disease relationships) [3]. The size constraint allows to forbid patterns that are too small w.r.t. the number of items (number of words) to be relevant patterns; we set ℓ_{min} to 3. The membership constraint enables to filter out sequential patterns that do not contain some selected items. For example, we state that extracted patterns must contain at least the two items GENE and DISEASE. We used the $gap[0, 9]$ constraint, which is the best setting found in [3]. As no specialized method exists for this combination of constraints, we thus compare GAP−SEQ with and without additional constraints.

Table 4 reports, for each value of $minsup$, the number of patterns extracted and the associated CPU times, the number of propagations and the number of nodes in the search tree. Additional constraints obviously restrict the number of extracted patterns. As the problem is more constrained, the size of the developed search tree is smaller. Even if the number of propagations is higher, the resulting CPU times are smaller. To conclude, thanks to the GAP−SEQ global constraint and its encoding, additional constraints like size, membership and regular expressions constraints can be easily stated.

(e) Evaluating the ability of GAP-SEQ to efficiently handle SPM. In order to simulate the absence of gap constraints, we used the ineffective $gap[0, \ell]$ constraint (recall that ℓ is the size of the longest sequence of SDB). We compared GAP−SEQ$[0, \ell]$ with PREFIX-PROJECTION and two configurations of cSpade for SPM: cSpade without gap constraint and cSpade with M and N set respectively to 0 and ℓ, denoted by cSpade$[0, \ell]$. Let us note that all the above methods will extract the same set of sequential patterns.

Fig. 5 reports the CPU times for the four methods. First, cSpade obtains the best performance (except on Protein). These results confirm those observed in [7].

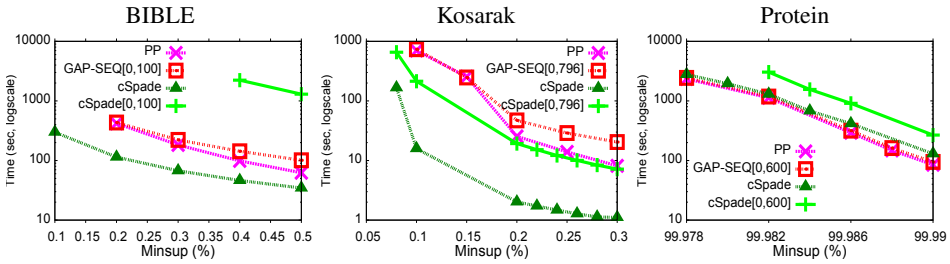

Fig. 5: Comparing GAP-SEQ with PREFIX-PROJECTION and cSpade for SPM on BIBLE, Kosarak and Protein.

Second, GAP−SEQ$[0, \ell]$ and PREFIX-PROJECTION exhibit similar behavior, even if GAP−SEQ$[0, \ell]$ is slightly less faster. So, even if GAP−SEQ handles both cases (with and without gap), it remains very competitive for SPM. Third, GAP−SEQ$[0, \ell]$ clearly outperforms cSpade$[0, \ell]$ (except on Kosarak). This is probably due to the huge number of unnecessary joining operations performed by cSpade$[0, \ell]$. To conclude, all the performed experiments demonstrate the ability of GAP-SEQ to efficiently handle SPM.

Finally, the gecode implementation of GAP−SEQ and the datasets used in our experiments are available online[7].

6 Conclusion

In this paper, we have introduced the global constraint GAP-SEQ enabling to handle SPM with or without gap constraints. The filtering algorithm takes benefits from the principle of right pattern extensions and prefix anti-monotonicity property of the gap constraint. GAP-SEQ enables to handle several types of constraints simultaneously and does not require any reified constraints nor any extra variables to encode the subsequence relation. Experiments performed on several real-life datasets (i) show that our approach clearly outperforms existing CP approaches as well as specialized methods for GSPM on large datasets, and (ii) demonstrate the ability of GAP-SEQ to efficiently handle SPM.

This work opens several issues for future researches. We plan to handle constraints on set of sequential patterns such as closedness, relevant subgroup and skypattern constraints.

References

1. Agrawal, R., Srikant, R.: Mining sequential patterns. In: Yu, P.S., Chen, A.L.P. (eds.) ICDE. pp. 3–14. IEEE Computer Society (1995)
2. Ayres, J., Flannick, J., Gehrke, J., Yiu, T.: Sequential pattern mining using a bitmap representation. In: KDD 2002. pp. 429–435. ACM (2002)
3. Béchet, N., Cellier, P., Charnois, T., Crémilleux, B.: Sequential pattern mining to discover relations between genes and rare diseases. In: CBMS (2012)

[7] https://sites.google.com/site/prefixprojection4cp/

4. Coquery, E., Jabbour, S., Saïs, L., Salhi, Y.: A SAT-based approach for discovering frequent, closed and maximal patterns in a sequence. In: ECAI. pp. 258–263 (2012)
5. Fournier-Viger, P., Gomariz, A., Gueniche, T., Soltani, A., Wu, C., Tseng, V.: SPMF: A Java Open-Source Pattern Mining Library. J. of Machine Learning Resea. 15, 3389–3393 (2014), http://jmlr.org/papers/v15/fournierviger14a.html
6. Ji, X., Bailey, J., Dong, G.: Mining minimal distinguishing subsequence patterns with gap constraints. In: (ICDM'05. pp. 194–201 (2005)
7. Kemmar, A., Loudni, S., Lebbah, Y., Boizumault, P., Charnois, T.: PREFIX-PROJECTION global constraint for sequential pattern mining. In: Principles and Practice of Constraint Programming - 21st International Conference, CP 2015, Cork, Ireland, August 31 - September 4, 2015, Proceedings. pp. 226–243 (2015)
8. Kemmar, A., Ugarte, W., Loudni, S., Charnois, T., Lebbah, Y., Boizumault, P., Crémilleux, B.: Mining relevant sequence patterns with cp-based framework. In: ICTAI. pp. 552–559 (2014)
9. Li, C., Wang, J.: Efficiently mining closed subsequences with gap constraints. In: SIAM 2008 on Data Mining. pp. 313–322 (2008)
10. Li, C., Yang, Q., Wang, J., Li, M.: Efficient mining of gap-constrained subsequences and its various applications. Trans. Knowl. Discov. Data 6(1), 2:1–2:39 (Mar 2012)
11. Métivier, J.P., Loudni, S., Charnois, T.: A constraint programming approach for mining sequential patterns in a sequence database. In: ECML/PKDD Workshop on Languages for Data Mining and Machine Learning (2013)
12. Négrevergne, B., Guns, T.: Constraint-based sequence mining using constraint programming. In: CPAIOR'15. pp. 288–305 (2015)
13. Pei, J., Han, J., Mortazavi-Asl, B., Pinto, H., Chen, Q., Dayal, U., Hsu, M.: PrefixSpan: Mining sequential patterns by prefix-projected growth. In: ICDE. pp. 215–224. IEEE Computer Society (2001)
14. Wu, X., Zhu, X., He, Y., Arslan, A.N.: PMBC: pattern mining from biological sequences with wildcard constraints. Comp. in Bio. and Med. 43(5), 481–492 (2013)
15. Yang, G.: Computational aspects of mining maximal frequent patterns. Theor. Comput. Sci. 362(1-3), 63–85 (2006)
16. Zaki, M.J.: Sequence mining in categorical domains: Incorporating constraints. In: CIKM'00. pp. 422–429 (2000)
17. Zaki, M.J.: SPADE: An efficient algorithm for mining frequent sequences. Machine Learning 42(1/2), 31–60 (2001)
18. Zhang, M., Kao, B., Cheung, D.W., Yip, K.Y.: Mining periodic patterns with gap requirement from sequences. TKDD 1(2) (2007)